

```
[root@nanyuki scripts]# sh 002-run-sqoop-import.sh
```

```
Warning: /usr/hdp/2.6.5.0-292/accumulo does not exist! Accumulo imports will fail.
```

```
Please set $ACCUMULO_HOME to the root of your Accumulo installation.
```

```
18/12/17 13:29:43 INFO sqoop.Sqoop: Running Sqoop version: 1.4.6.2.6.5.0-292
```

```
Enter password:
```

```
18/12/17 13:29:48 INFO tool.BaseSqoopTool: Using Hive-specific delimiters for output. You can override
```

```
18/12/17 13:29:48 INFO tool.BaseSqoopTool: delimiters with --fields-terminated-by, etc.
```

```
18/12/17 13:29:48 INFO manager.MySQLManager: Argument '--fetch-size 1' will probably get ignored by MySQL JDBC driver.
```

```
18/12/17 13:29:48 INFO tool.CodeGenTool: Beginning code generation
```

```
18/12/17 13:29:51 INFO manager.SqlManager: Executing SQL statement: SELECT t.* FROM `test_lineage` AS t LIMIT 1
```

```
18/12/17 13:29:51 INFO manager.SqlManager: Executing SQL statement: SELECT t.* FROM `test_lineage` AS t LIMIT 1
```

```
18/12/17 13:29:51 INFO orm.CompilationManager: HADOOP_MAPRED_HOME is /usr/hdp/2.6.5.0-292/hadoop-mapreduce
```

```
Note: /tmp/sqoop-root/compile/e487a80e27a3e93e701d566d32419422/test_lineage.java uses or overrides a deprecated API.
```

```
Note: Recompile with -Xlint:deprecation for details.
```

```
18/12/17 13:30:24 INFO orm.CompilationManager: Writing jar file: /tmp/sqoop-root/compile/e487a80e27a3e93e701d566d32419422/test_lineage.jar
```

```
18/12/17 13:30:24 WARN manager.MySQLManager: It looks like you are importing from mysql.
```

```
18/12/17 13:30:24 WARN manager.MySQLManager: This transfer can be faster! Use the --direct
```

```
18/12/17 13:30:24 WARN manager.MySQLManager: option to exercise a MySQL-specific fast path.
```

```
18/12/17 13:30:24 INFO manager.MySQLManager: Setting zero DATETIME behavior to convertToNull (mysql)
```

```
18/12/17 13:30:25 INFO mapreduce.ImportJobBase: Beginning import of test_lineage
```

```
18/12/17 13:30:52 INFO client.RMProxy: Connecting to ResourceManager at nanyuki.kenya.ke/192.168.0.31:8050
```

```
18/12/17 13:30:57 INFO client.AHSProxy: Connecting to Application History server at nanyuki.kenya.ke/192.168.0.31:10200
```

```
18/12/17 13:32:38 INFO db.DBInputFormat: Using read committed transaction isolation
```

```
18/12/17 13:32:39 INFO mapreduce.JobSubmitter: number of splits:1
```

```
18/12/17 13:32:40 INFO mapreduce.JobSubmitter: Submitting tokens for job: job_1545046661241_0003
```

```
18/12/17 13:32:42 INFO impl.YarnClientImpl: Submitted application application_1545046661241_0003
```

```
18/12/17 13:32:43 INFO mapreduce.Job: The url to track the job: http://nanyuki.kenya.ke:8088/proxy/application_1545046661241_0003/
```

```
18/12/17 13:32:43 INFO mapreduce.Job: Running job: job_1545046661241_0003
```

```
18/12/17 13:35:52 INFO mapreduce.Job: Job job_1545046661241_0003 running in uber mode : false
```

```
18/12/17 13:35:52 INFO mapreduce.Job: map 0% reduce 0%
```

```
18/12/17 13:36:52 INFO mapreduce.Job: map 100% reduce 0%
```

```
18/12/17 13:37:01 INFO mapreduce.Job: Job job_1545046661241_0003 completed successfully
```

```
18/12/17 13:37:02 INFO mapreduce.Job: Counters: 30
```

#### File System Counters

```
FILE: Number of bytes read=0
```

```
FILE: Number of bytes written=168945
```

```
FILE: Number of read operations=0
```

```
FILE: Number of large read operations=0
```

```
FILE: Number of write operations=0
```

```
HDFS: Number of bytes read=87
```

```
HDFS: Number of bytes written=0
```

```
HDFS: Number of read operations=4
```

```
HDFS: Number of large read operations=0
```

```
HDFS: Number of write operations=2
```

#### Job Counters

```
Launched map tasks=1
```

```
Other local map tasks=1
```

```
Total time spent by all maps in occupied slots (ms)=111020
```

```
Total time spent by all reduces in occupied slots (ms)=0
```

```
Total time spent by all map tasks (ms)=55510
```

```
Total vcore-milliseconds taken by all map tasks=55510
```

```
Total megabyte-milliseconds taken by all map tasks=85263360
```

#### Map-Reduce Framework

```
Map input records=0
```

```
Map output records=0
```

```
Input split bytes=87
```

```
Spilled Records=0
```

```
Failed Shuffles=0
```

```
Merged Map outputs=0
```

```
GC time elapsed (ms)=81
```

```
CPU time spent (ms)=2470
```

```
Physical memory (bytes) snapshot=178155520
```

```
Virtual memory (bytes) snapshot=3303743488
```

```
Total committed heap usage (bytes)=89653248
```

#### File Input Format Counters

```
Bytes Read=0
```

#### File Output Format Counters

```
Bytes Written=0
```

```
18/12/17 13:37:02 INFO mapreduce.ImportJobBase: Transferred 0 bytes in 374.8356 seconds (0 bytes/sec)
```

```
18/12/17 13:37:02 INFO mapreduce.ImportJobBase: Retrieved 0 records.
```

18/12/17 13:37:02 INFO mapreduce.ImportJobBase: Publishing Hive/Hcat import job data to Listeners  
18/12/17 13:37:05 INFO atlas.ApplicationProperties: Looking for atlas-application.properties in classpath  
18/12/17 13:37:05 INFO atlas.ApplicationProperties: Loading atlas-application.properties from file:/etc/sqoop/2.6.5.0-292/0/atlas-application.properties  
18/12/17 13:37:06 ERROR security.InMemoryJAASConfiguration: Unable to add JAAS configuration for client [KafkaClient] as it is missing param [atlas.jaas.KafkaClient.loginModuleName]. Skipping JAAS config for [KafkaClient]  
18/12/17 13:37:08 INFO kafka.KafkaNotification: ==> KafkaNotification()  
18/12/17 13:37:08 INFO kafka.KafkaNotification: <== KafkaNotification()  
18/12/17 13:37:08 INFO hook.AtlasHook: Created Atlas Hook  
18/12/17 13:37:12 INFO kafka.KafkaNotification: ==> KafkaNotification.createProducer()  
18/12/17 13:37:12 INFO producer.ProducerConfig: ProducerConfig values:  
acks = 1  
batch.size = 16384  
bootstrap.servers = [nanyuki.kenya.ke:6667]  
buffer.memory = 33554432  
client.id =  
compression.type = none  
connections.max.idle.ms = 540000  
enable.idempotence = false  
interceptor.classes = null  
key.serializer = class org.apache.kafka.common.serialization.StringSerializer  
linger.ms = 0  
max.block.ms = 60000  
max.in.flight.requests.per.connection = 5  
max.request.size = 1048576  
metadata.max.age.ms = 300000  
metric.reporters = []  
metrics.num.samples = 2  
metrics.recording.level = INFO  
metrics.sample.window.ms = 30000  
partitioner.class = class org.apache.kafka.clients.producer.internals.DefaultPartitioner  
receive.buffer.bytes = 32768  
reconnect.backoff.max.ms = 1000  
reconnect.backoff.ms = 50  
request.timeout.ms = 30000  
retries = 0  
retry.backoff.ms = 100  
sasl.jaas.config = null  
sasl.kerberos.kinit.cmd = /usr/bin/kinit  
sasl.kerberos.min.time.before.relogin = 60000  
sasl.kerberos.service.name = null  
sasl.kerberos.ticket.renew.jitter = 0.05  
sasl.kerberos.ticket.renew.window.factor = 0.8  
sasl.mechanism = GSSAPI  
security.protocol = PLAINTEXT  
send.buffer.bytes = 131072  
ssl.cipher.suites = null  
ssl.enabled.protocols = [TLSv1.2, TLSv1.1, TLSv1]  
ssl.endpoint.identification.algorithm = null  
ssl.key.password = null  
ssl.keymanager.algorithm = SunX509  
ssl.keystore.location = null  
ssl.keystore.password = null  
ssl.keystore.type = JKS  
ssl.protocol = TLS  
ssl.provider = null  
ssl.secure.random.implementation = null  
ssl.trustmanager.algorithm = PKIX  
ssl.truststore.location = null  
ssl.truststore.password = null  
ssl.truststore.type = JKS  
transaction.timeout.ms = 60000  
transactional.id = null  
value.serializer = class org.apache.kafka.common.serialization.StringSerializer  
18/12/17 13:37:18 WARN producer.ProducerConfig: The configuration 'key.deserializer' was supplied but isn't a known config.  
18/12/17 13:37:18 WARN producer.ProducerConfig: The configuration 'value.deserializer' was supplied but isn't a known config.  
18/12/17 13:37:18 WARN producer.ProducerConfig: The configuration 'hook.group.id' was supplied but isn't a known config.  
18/12/17 13:37:18 WARN producer.ProducerConfig: The configuration 'zookeeper.connection.timeout.ms' was supplied but isn't a known config.  
18/12/17 13:37:18 WARN producer.ProducerConfig: The configuration 'zookeeper.session.timeout.ms' was supplied but isn't a known config.

18/12/17 13:37:18 WARN producer.ProducerConfig: The configuration 'enable.auto.commit' was supplied but isn't a known config.  
18/12/17 13:37:18 WARN producer.ProducerConfig: The configuration 'zookeeper.connect' was supplied but isn't a known config.  
18/12/17 13:37:18 WARN producer.ProducerConfig: The configuration 'zookeeper.sync.time.ms' was supplied but isn't a known config.  
18/12/17 13:37:18 WARN producer.ProducerConfig: The configuration 'session.timeout.ms' was supplied but isn't a known config.  
18/12/17 13:37:18 WARN producer.ProducerConfig: The configuration 'auto.offset.reset' was supplied but isn't a known config.  
18/12/17 13:37:18 INFO utils.AppInfoParser: Kafka version : 1.0.0.2.6.5.0-292  
18/12/17 13:37:18 INFO utils.AppInfoParser: Kafka commitId : 2ff1ddae17fb8503  
18/12/17 13:37:18 INFO kafka.KafkaNotification: <== KafkaNotification.createProducer()  
18/12/17 13:37:20 INFO manager.SqlManager: Executing SQL statement: SELECT t.\* FROM `test\_lineage` AS t LIMIT 1  
18/12/17 13:37:20 INFO hive.HiveImport: Loading uploaded data into Hive  
Logging initialized using configuration in jar:file:/usr/hdp/2.6.5.0-292/hive/lib/hive-common-1.2.1000.2.6.5.0-292.jar!/hive-log4j.properties  
OK  
Time taken: 6.966 seconds  
Loading data to table default.test\_sqoop\_lineage  
Table default.test\_sqoop\_lineage stats: [numFiles=2, numRows=0, totalSize=0, rawDataSize=0]  
OK  
Time taken: 5.352 seconds